Carnegie Mellon University
Computer Science Department

THESIS ORAL

BRUNO VAVALA

Secure Large-Scale Outsourced Services Founded on Trustworthy Code Executions

Monday, July 17, 2017 – 11:30 a.m. – NSH 1507

The Cloud Computing model has incentivized companies to outsource services to third-party providers. Service owners can use third-party computational, storage and network resources while avoiding the cost of acquiring an IT infrastructure. However, they have to rely on the trustworthiness of the third-party providers, who ultimately need to guarantee that the services run as intended.

This thesis shows how to secure the execution of large-scale services. From the perspective of a client that sends a request and receives a response, trust can be established by verifying a small proof of correct execution that is attached to the result. On the remote provider's platform, a small trusted computing base enables the secure execution of generic services composed of a large source code and/or working on large data sets, using an abstraction layer that is implementable on diverse trusted hardware architectures.

Our small TCB implements three orthogonal techniques that are the core contributions of this thesis. The first one targets the identification (and the execution) of only the part of code that is necessary to fulfill a client's request. This allows an increase both in security and efficiency by leaving any code that is not required to run the service outside the execution environment. The second contribution enables terabyte-scale data processing by means of a secure in-memory data handling mechanism. This allows a service to retrieve data that is validated on access and before use. Notably, data I/O is performed using virtual memory mechanisms that do not require any system call from the trusted execution environment, thereby reducing the attack surface. The third contribution is a novel fully-passive secure replication scheme that is tolerant to software attacks. Fault-tolerance delivers availability guarantees to clients, while passive replication allows for computationally efficient processing. Interestingly, all of our techniques are based on the same abstraction layer of the trusted hardware. In addition, our implementation and experimental evaluation demonstrate the practicality of these approaches.

Thesis Committee:
Peter Steenkiste, Co-Chair
Nuno Neves, Co-Chair
Anupam Datta
Vyas Sekar
Antonia Lopes, University of Lisbon